arXiv:1605.06122v1 [stat.CO] 19 May 2016

MCMC with Strings and Branes: The Suburban Algorithm

Jonathan J. Heckman,[1] Jeffrey G. Bernstein,[2] and Ben Vigoda[3]

[1]*Department of Physics, University of North Carolina at Chapel Hill, NC*
[2]*Analog Devices | Lyric Labs, Cambridge, MA*
[3]*Gamalon Labs, Cambridge, MA*

Motivated by the physics of strings and branes, we introduce a general suite of Markov chain Monte Carlo (MCMC) "suburban samplers" (i.e., spread out Metropolis). The suburban algorithm involves an ensemble of statistical agents connected together by a random network. Performance of the collective in reaching a fast and accurate inference depends primarily on the average number of nearest neighbor connections. Increasing the average number of neighbors above zero initially leads to an increase in performance, though there is a critical connectivity with effective dimension $d_{\rm eff} \sim 1$, above which "groupthink" takes over, and the performance of the sampler declines.

I. INTRODUCTION

Markov chain Monte Carlo (MCMC) methods are a remarkably robust way to sample from complex probability distributions. Metropolis-Hastings (MH) sampling [1, 2] stands out as an important benchmark. An appealing feature of MH sampling is the simple physical picture which underlies the general method. Roughly speaking, the idea is that the thermal fluctuations of a particle moving in an energy landscape provides a conceptually elegant way to sample from a target distribution.

But there are also potential drawbacks to MCMC methods. For example, the speed of convergence to the correct posterior is often unknown since a sampler can become trapped in a metastable equilibrium for a long period of time. Once a sampler becomes trapped, a large free energy barrier can obstruct an accurate determination of the distribution. From this perspective it is therefore natural to ask whether further inspiration from physics can lead to new examples of samplers.

Now, although the physics of point particles underlies much of our modern understanding of natural phenomena, it has proven fruitful to consider objects such as strings and branes with finite extent in p spatial dimensions (a string being a case of a 1-brane). One of the main features of branes is that the number of spatial dimensions strongly affects how a localized perturbation propagates across its worldvolume. Viewing a brane as a collective of point particles that interact with one another (see fig. 1), this suggests applications to questions in statistical inference [3].

Motivated by these physical considerations, our aim in this work will be to study generalizations of the MH algorithm for such extended objects. For an ensemble of M parallel MH samplers of a distribution $\pi(x)$, we can alternatively view this as a single particle sampling from M variables $x_1, ..., x_M$ with density:

$$\pi(x_1, ..., x_M) = \pi(x_1)...\pi(x_M), \tag{1}$$

where the proposal kernel is simply:

$$q_{\rm par}(x_1^{\rm new}, ..., x_M^{\rm new}|x_1^{\rm old}, ..., x_M^{\rm old}) = \prod_{\sigma=1}^{M} q(x_\sigma^{\rm new}|x_\sigma^{\rm old}). \tag{2}$$

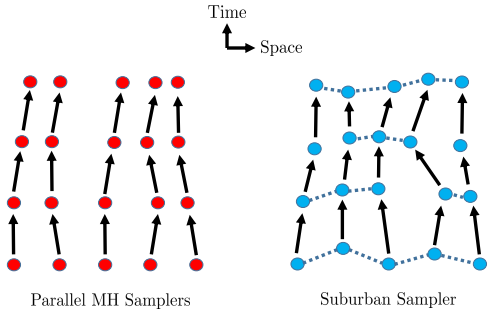

FIG. 1: Depiction of how parallel MH samplers (left) and a suburban sampler (right) evolve as a function of time.

To realize MCMC with strings and branes, we keep the same target $\pi(x_1, ..., x_M)$, but we change the proposal kernel by interpreting the index σ on x_σ as specifying the location of a statistical agent in a network. Depending on the connectivity of this network, an agent may interact with several neighboring agents $\mathrm{Nb}(x_\sigma)$, so we introduce a proposal kernel:[1]

$$q_{\rm brane}(x_1^{\rm new}, ..., x_M^{\rm new}|x_1^{\rm old}, ..., x_M^{\rm old}) = \prod_{\sigma=1}^{M} q_\sigma(x_\sigma^{\rm new}|\mathrm{Nb}(x_\sigma^{\rm old})). \tag{3}$$

In the above, the connectivity of the extended object specifies its overall topology. For example, in the case of a string, i.e., a one-dimensional extended object, the neighbors of x_i are x_{i-1}, x_i, and x_{i+1}. Fig. 1 depicts the time evolution of parallel MH samplers compared with the suburban sampler.

However, there are potentially many consistent ways to connect together the inferences of statistical agents. From the perspective of physics, this amounts to a notion of distance/proximity between nearest neighbors in a brane. A physically well-motivated way to eliminate

[1] From this perspective, the suburban algorithm is a particular choice of ensemble MCMC (see e.g. [4–9]).

this arbitrary feature is to allow the notion of proximity itself to *fluctuate*, and for the brane to split and join.

We view MCMC with extended strings and branes as a novel class of ensemble samplers. By correlating the inferences of nearest neighbors, we can expect there to be some impact on performance. For example, the degree of connectivity impacts the mixing rate for obtaining independent samples. Another important feature is that because we are dealing with an extended object, different statistical agents may become localized in different high density regions. Provided the connectivity with neighbors is sufficiently low, coupling these agents then has the potential to provide a more accurate global characterization of a target distribution. Conversely, connecting too many agents together may cause the entire collective to suffer from "groupthink" in the sense of [3]. In particular, we shall present evidence that the optimal connectivity for a network of agents on a grid arranged as a hypercubic lattice with some percolation (i.e., we allow for broken links) occurs at a critical effective dimension:

$$d_{\text{eff}} \sim 1 \tag{4}$$

where $2d_{\text{eff}}$ is the average number of neighbors.

To summarize: With too few friends one drifts into oblivion, but with too many friends one becomes a boring conformist.

Turning our discussion around, one can view this paper as providing a concrete way to study the physics of branes with a strongly fluctuating worldvolume, that is, the nonperturbative regime of string theory.

The Appendices provide some additional details (see also [10]). The suburban code is available at `https://gitlab.com/suburban/suburban`.

II. MCMC WITH STRINGS AND BRANES

One of the main ideas we shall develop in this paper is MCMC methods for extended objects. In an MCMC algorithm we produce a sequence of "timesteps" $x^{(1)}, ..., x^{(t)}, ..., x^{(N)}$ which can be viewed as the motion of a point particle exploring a target space Ω. More formally, this sequence of points defines the "worldline" for a particle, and consequently a map:

$$t \mapsto x(t). \tag{5}$$

For an extended object with d spatial directions, we get a map from a "worldvolume" to the target:

$$(t, \sigma_1, ..., \sigma_d) \mapsto x(t, \sigma_1, ..., \sigma_d). \tag{6}$$

The special cases $d = 0$ and $d = 1$ respectively denote a point particle and string.

The general physical intuition is that minus the log of the target distribution $\pi(x) = \exp(-V(x))$ defines a potential energy, and minus the log of the Markov chain transition probability $T(x \to x') = \exp(-K(x, x'))$ is a kinetic energy. The key point is that statistical field theory in $d + 1$ Euclidean dimensions strongly depends on the number of dimensions. For example, the two-point function for a free Gaussian field $x(\sigma)$ with $\sigma \in \mathbb{R}^{d+1}$ is:

$$\mathbb{E}(x(\sigma)x(0)) \sim 1/||\sigma||^{d-1} \tag{7}$$

where in the case of $d = 1$, the two-point function is $\log ||\sigma||$. For $d \leq 1$, a random field explores its surroundings at large σ, but the overall variance decreases as $d \to 1$. For $d > 1$, however, "groupthink" sets in and the ensemble less quickly explores its surroundings. This suggests a special role for stringlike objects [3].

In a theory of quantum gravity (such as string theory) it is also physically natural to let the proximity of nearest neighbors fluctuate. So, we introduce an ensemble of random graphs $\mathcal{A}$. For example, for a d-dimensional toroidal hypercubic lattice, introduce m lattice sites along a spatial direction so that $m^d = M$ is the total number of agents. For a hypercubic lattice in d dimensions, we define the ensemble of random graphs for a brane $\mathcal{A}_{\text{brane}}(p_{\text{join}})$ as one in which we have a random shuffling of the agents, and in which a given link in a d-dimensional hypercubic lattice is active with probability p_{join}. We can also consider more general ensembles of adjacency matrices. For example, the Erdös-Renyi ensemble $\mathcal{A}_{\text{ER}}(p_{\text{join}})$ has an edge between any two nodes with probability p_{join}. We also introduce the notion of an effective dimension which depends on the average number of neighbors:

$$d_{\text{eff}} = n_{\text{avg}}/2, \tag{8}$$

which need not be an integer.

III. THE SUBURBAN ALGORITHM

We now present the suburban algorithm. For ease of exposition, we shall present the case of a 1D target. The generalization to a D-dimensional target is straightforward, and we can take MH within a Gibbs sampler, or a sampler with joint variables in which all D dimensions update simultaneously.

To avoid overloading the notation, we shall write $\mathcal{X}^{(t)} \equiv \left\{x_1^{(t)}, ..., x_M^{(t)}\right\}$ for the current state of the grid. Instead of directly sampling from $\pi(x)$, we introduce multiple copies of the target and sample from the joint distribution $\pi(\mathcal{X}) = \pi(x_1)...\pi(x_M)$ using MH sampling with proposal kernel $q_{\text{brane}}(x_1^{\text{new}}, ..., x_M^{\text{new}}|x_1^{\text{old}}, ..., x_M^{\text{old}}) = q(\mathcal{X}^{(\text{new})}|\mathcal{X}^{(\text{old})}, A)$, where A denotes the adjacency matrix. If the system is in a state $\mathcal{X}^{(t)}$, with adjacency matrix $A^{(t)}$, we pick a new state according to the MH update rule with a proposal kernel which depends on both these inputs. The MH acceptance probability is:

$$a\left(\mathcal{X}^{\text{new}}|\mathcal{X}^{\text{old}}, A\right) = \min\left(1, \frac{q(\mathcal{X}^{\text{old}}|\mathcal{X}^{\text{new}}, A)}{q(\mathcal{X}^{\text{new}}|\mathcal{X}^{\text{old}}, A)} \frac{\pi\left(\mathcal{X}^{\text{new}}\right)}{\pi\left(\mathcal{X}^{\text{old}}\right)}\right) \tag{9}$$

This leads us to algorithm 1.

Some of these steps can be parallelized whilst retaining detailed balance. For example we could pick a coloring

Algorithm 1 Suburban Sampler

Randomly Initialize $\mathcal{X}^{(0)}$ and $A^{(0)}$
for $t = 0$ to $N-1$ **do**
  $\mathcal{X}^{(*)} \leftarrow$ sample from $q(\mathcal{X}|\mathcal{X}^{(t)}, A^{(t)})$
  accept with probability $a(\mathcal{X}^{*}|\mathcal{X}^{(t)}, A^{(t)})$
    if accept = true **then**
      $\mathcal{X}^{(t+1)} \leftarrow \mathcal{X}^{(*)}$
    else
      $\mathcal{X}^{(t+1)} \leftarrow \mathcal{X}^{(t)}$
  $A^{(t+1)} \leftarrow$ draw from $\mathcal{A}$
return $\mathcal{X}^{(1)}, ..., \mathcal{X}^{(N)}$

of a graph and then perform an update for all nodes of a particular color whilst holding fixed the rest. We can also stochastically evolve the adjacency matrices.

Now, having collected a sequence of values $\mathcal{X}^{(1)}, ..., \mathcal{X}^{(N)}$, we can interpret this as $N \times M$ samples of the original distribution $\pi(x)$. As standard for MCMC methods, we can then calculate quantities of interest such as the mean:

$$\overline{x} \simeq \frac{1}{NM} \sum_{\sigma,t} x_\sigma^{(t)} \tag{10}$$

as well as higher order moments.

Let us discuss the reason we expect our sampler to converge to the correct posterior distribution. First note that although we are modifying the proposal kernel at each time step (i.e., by introducing a different adjacency matrix $A \in \mathcal{A}$), this modification is independent of the current state of the system. So, it cannot impact the eventual posterior distribution we obtain. Second, we observe that since we are just performing a specific kind of MH sampling routine for the distribution $\pi(x_1, ..., x_M)$, we expect to converge to the correct posterior distribution. But, since the variables $x_1, ..., x_M$ are all independent, this is tantamount to having also sampled multiple times from $\pi(x)$. The caveat is that we need the sampler to actually wander around during its random walk; $d \leq 1$ is typically necessary to prevent "groupthink."

A. Implementation

To accommodate a flexible framework for prototyping, we have implemented the suburban algorithm in the probabilistic programming language `Dimple` [11].

For practical purposes we take a fairly large burn-in cut, discarding the first 10% of samples from a run. We always perform Gibbs sampling over the M agents. For MH within Gibbs sampling over a D-dimensional target, we thus get a Gibbs schedule with $D \times M$ updates for each time step. For a joint sampler, the Gibbs schedule consists of just M updates.

The specific choice of $q_\sigma(x_\sigma|\mathrm{Nb}(x_\sigma))$ for eqn. (3) is motivated by having a free Gaussian field on a fluctuating graph topology:

$$\propto \exp\left(-\alpha_\sigma \left(D_t x_\sigma\right)^2 - \sum_{n(\sigma)} \beta \left(D_t x_\sigma - D_{n(\sigma)} x_\sigma\right)^2\right) \tag{11}$$

with:

$$D_t x_\sigma = x_\sigma^{(t+1)} - x_\sigma^{(t)} \tag{12}$$

$$D_{n(\sigma)} x_\sigma = x_{n(\sigma)}^{(t)} - x_\sigma^{(t)} \tag{13}$$

for $n(\sigma)$ a neighbor of σ on the graph defined by the adjacency matrix. Additionally, we set the hyperparameters for the kernel as:

$$\alpha_\sigma = 2\beta - n_\sigma^{\mathrm{tot}} \beta, \tag{14}$$

that is, we take an adaptive value for α_σ specified by the number of nearest neighbors joined to x_σ. This condition leads to a well-behaved continuum limit on a fully connected hypercubic lattice.

IV. NUMERICAL EXPERIMENTS

In most cases, we consider MH within Gibbs sampling, though we also consider the case where joint variables are sampled, that is, pure MH. Rather than perform error analysis within a single long MCMC run, we opt to take multiple independent trials of each MCMC run in which we vary the hyperparameters of the sampler such as the overall topology and average degree of connectivity of the sampler. Though this leads to less efficient statistical estimators, it has the virtue of allowing us to easily compare the performance of different algorithms, i.e., as we vary the continuous and discrete hyperparameters of the suburban algorithm. We take $M = 81 = 9^2 = 3^4$ to compare different grid topologies. We have also compared performance with parallel slice (within Gibbs) samplers [12] to ensure that our performance is comparable to other benchmarks.

To gauge accuracy, we collect the inferred mean and covariance matrix. We then compute the distance to the true values:

$$d_{\mathrm{mean}} \equiv \|\mu_{\mathrm{inf}} - \mu_{\mathrm{true}}\| \tag{15}$$

$$d_{\mathrm{cov}} \equiv \left(\mathrm{Tr}\left(\left(\Sigma_{\mathrm{inf}} - \Sigma_{\mathrm{true}}\right) \cdot \left(\Sigma_{\mathrm{inf}} - \Sigma_{\mathrm{true}}\right)^T\right)\right)^{1/2}. \tag{16}$$

We also collect performance metrics from the MCMC runs such as the rejection rate. A typical rule of thumb is that for targets with no large free energy barriers, a rejection rate of somewhere between $50\% - 80\%$ is acceptable (see e.g., [13]). We also collect the integrated auto-correlation time for the "energy" of the distribution:

$$V = -\log \pi(x_1, ..., x_M), \tag{17}$$

by collecting the values $V^{(1)}, ..., V^{(N)}$. For $-N < k < N$, we evaluate:

$$c(k) \equiv \left\{ \begin{array}{ll} \frac{1}{N} \sum_{t=1}^{N-k} \left(V^{(t)} - \overline{V}\right) \left(V^{(t+k)} - \overline{V}\right) & k \geq 0 \\ \frac{1}{N} \sum_{t=1}^{N+k} \left(V^{(t)} - \overline{V}\right) \left(V^{(t-k)} - \overline{V}\right) & k < 0 \end{array} \right\}, \tag{18}$$

and then extract the integrated auto-correlation time:

$$\tau_{\rm dec} \equiv \sum_{-N<k<N} \left(1 - \left|\frac{k}{N}\right|\right) \left|\frac{c(k)}{c(0)}\right| \tag{19}$$

we also refer to this as the "decay time"as it reflects how quickly the chain mixes. For this observable we include all samples (no burn-in).

To extract numerical estimates we perform T independent trials with random initialization for each agent on $[-100, +100]^D$. We present all plots with a 3-sigma level standard error around the mean value from these trials. In practice, we typically find acceptable error bars for $T = 100$ and $T = 1000$ trials.

A. Effective Connectivity

Perhaps the single most important feature of the suburban algorithm is that it correlates the inferences drawn by nearest neighbors on a grid. Quite strikingly, we find that the effective dimension rather than the overall topology of the grid plays the dominant role in the performance of the algorithm.

We illustrate this point with a class of target distribution examples which we refer to as "symmetric mixtures."For a fixed choice of D the number of target space dimensions, we introduce a mixture model consisting of $2D$ equal weight components, each of which is a normal distribution with means and covariance matrices:

$$\mu_i^{(\pm,j)} = \pm\mu \times \delta_i^j \quad \Sigma^{(\pm,i)} = \sigma^2 \times \mathbb{I}_{D\times D}, \tag{20}$$

where $i, j = 1, ..., D$, δ_i^j is a Kronecker delta and $\mathbb{I}_{D\times D}$ is the $D \times D$ identity matrix. We find qualitatively similar behavior for $D = 2$ and $D = 10$, so we give the plots for the $D = 2$ runs with $\mu = 1.5$ and $\sigma^2 = 0.25$. In this case, we have four equally weighted components of our mixture model, and there is a free energy barrier separating these centers.

As a first class of tests, we consider sampling with different topology grids for $N = 10,000$ timesteps, with each grid consisting of $M = 81$ agents. For each choice of hyperparameter, we perform $T = 100$ trials.

We have scanned the value of β in steps of factors of 10, and find that performance is better around $\beta = 0.01$, so we focus on this case. In all cases, we find that the values of the observables $d_{\rm mean}$ and $d_{\rm cov}$ are comparable and small, indicating reasonable convergence.

There is, however, a marked difference in the mixing rate as we vary the split / join probability for the ensemble. In fig. 2 we display the values of $\tau_{\rm dec}$ as a function of the effective dimension dictated by the split / join rate for a given grid topology. Quite striking is the universal behavior of the samplers as a function of the effective dimension near $d_{\rm eff} \sim 1$, i.e., for connectivity similar to that of a string. Near $d_{\rm eff} = 0$, i.e., for parallel MH samplers, we also see much slower mixing rates. Once we go beyond $d_{\rm eff} \gtrsim 1$, the overall performance of the sampler suffers.

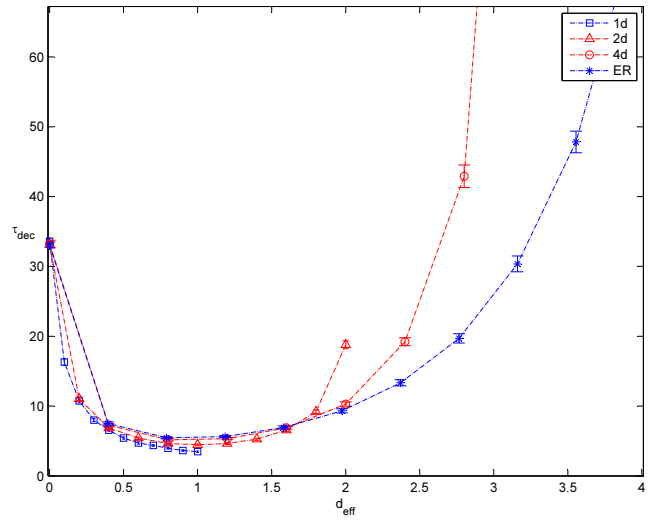

FIG. 2: Plot of the mixing rate $\tau_{\rm dec}$ as a function of the effective dimension of the grid. All data comes from sampling the $D = 2$ symmetric mixture model.

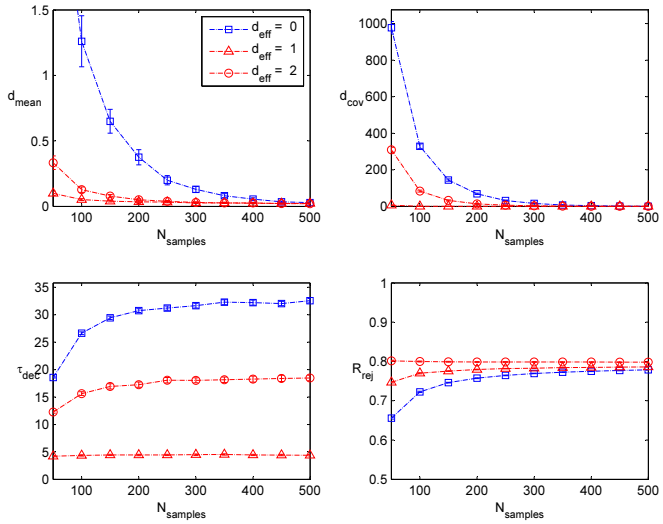

FIG. 3: Plots of the 2D symmetric mixture model tests.

Because of this universal behavior, we shall primarily focus on "representative behavior" as obtained from a $2d$ grid topology. In fig. 3 we show various performance metrics as a function of the total number of samples. By inspection, stringlike samplers tend to fare the best.

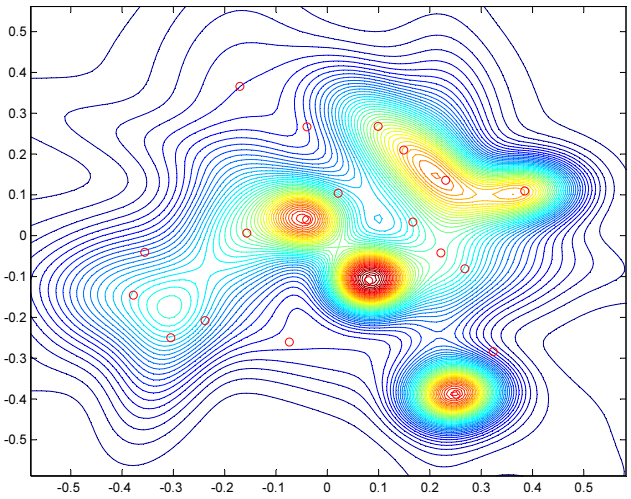

FIG. 4: Contour plot of the random mixture model with twenty components and random seed of 40. Red circles denote centers of individual components.

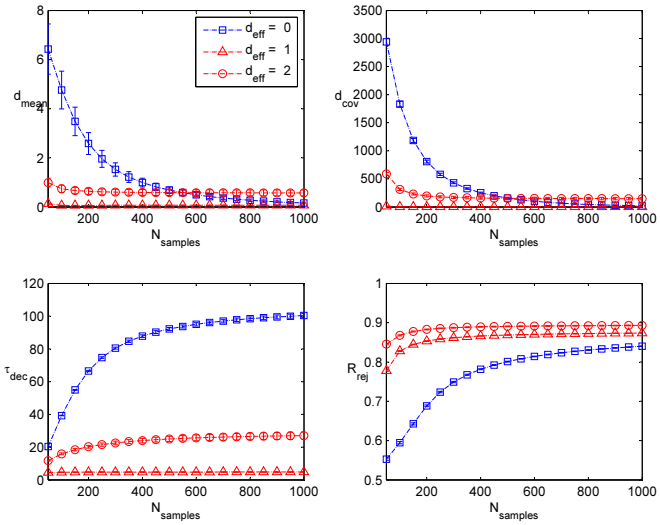

FIG. 5: Plots of the random seed 40 landscape tests.

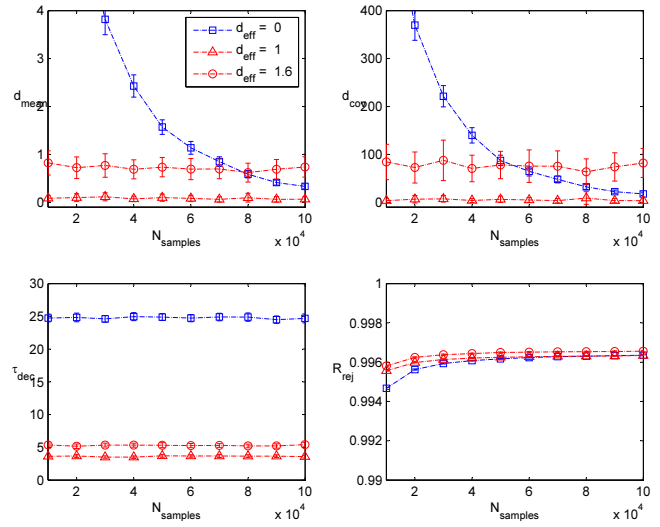

FIG. 6: Plots of the banana distribution tests.

B. Random Landscapes

As another example, we consider mixture models in which there is a landscape of local maxima and minima. A priori, a compromise will need to be struck between "wandering freely" and moving more slowly around individual components of the mixture model.

We use a variant on the same random mixture model considered in [14], focussing on the case of 20 Gaussian mixtures with relative weights randomly drawn from the uniform distribution on $[0, 1]$. Compared with [14], we take the parameters `stdmu` $= 0.4$, `stdsig` $= 10.0$. We do this primarily to achieve convergence for the samplers in a reasonable amount of time. The different mixture models are obtained by setting the random seed in the code of [14] to different values (see fig. 4).

By design, we have chosen our domain for the random variables so that the brane tension $\beta = 0.01$ should give a roughly comparable class of length scales for the target distribution. Since the overall topology of the grid does not appear to affect the qualitative behavior of the sampler, we have also focussed on the case of a $2d$ grid topology with shuffling and percolation. For each choice of hyperparameter, we perform $T = 100$ trials.

The random seed 40 model gives representative behavior. The stringlike sampler is faster and more accurate than the parallel MH sampler, while a $d_{\text{eff}} \sim 2$ sampler suffers from "groupthink," settling in an incorrect metastable configuration.

C. Banana Distribution

It is also of interest to consider distributions concentrated on a lower-dimensional subspace such as the two-dimensional "banana distribution":

$$\pi_{\text{banana}}(x, y) = \frac{1}{10\pi} \exp\left(-(x-1)^2 - 100(y-x^2)^2\right). \tag{21}$$

This distribution is often used as a performance test of various optimization algorithms.

We focus on a suburban sampler with joint variables, taking different grid topologies for the statistical agents and then perform a sweep over different values of the hyperparemeters β and d_{eff}, performing $T = 100$ trials for each case.

We present the representative case of a $2d$ grid, and further specialize to the tuned case of $\beta = 0.01$. Fig. 6 shows that parallel samplers ($d_{eff} = 0$) and collectives with groupthink ($d_{eff} > 1$) both fare worse than a stringlike sampler.

D. Free Energy Barriers

The extended nature of the suburban sampler also suggests that for target distributions with various disconnected "deep pockets," different pieces of the ensemble can wander over to different regions. We consider a mixture model with two Gaussian components:

$$\pi_{\rm GMM}(x) = \frac{3}{4}\mathcal{N}(x|\mu^{(+)},\Sigma) + \frac{1}{4}\mathcal{N}(x|\mu^{(-)},\Sigma), \qquad (22)$$

with:

$$\mu^{(\pm)} = (\pm L_{\rm barrier}, 0, ..., 0) \quad \Sigma = \sigma^2 \times \mathbb{I}_{D\times D}, \qquad (23)$$

where we vary $L_{\rm barrier}$ and hold fixed $\sigma = 0.25$. We take $N_{\rm samples} = 1000$ samples with $M = 81$ agents on a $2d$ grid with $\beta = 0.01$, performing $T = 1000$ independent trials. Since we use MH within Gibbs, we do not find much decrease in performance in comparing the $D = 2$ and $D = 10$ free energy barrier tests.

Fig. 7 shows that parallel samplers fare worse than the extended objects. The $d_{\rm eff} = 2$ runs are sometimes more accurate, but mix slower than for $d_{\rm eff} = 1$. After thinning samples, the former runs will be less accurate.

Acknowledgements

JJH thanks D. Krohn for collaboration at an early stage. We thank J.A. Barandes, C. Barber, C. Freer, M. Freytsis, J.J. Heckman Sr., A. Murugan, P. Oreto, R. Yang, and J. Yedidia for helpful discussions. The work of JJH is supported by NSF CAREER grant PHY-1452037. JJH also acknowledges support from the Bahnson Fund as well as the R. J. Reynolds Industries, Inc. Junior Faculty Development Award at UNC Chapel Hill.

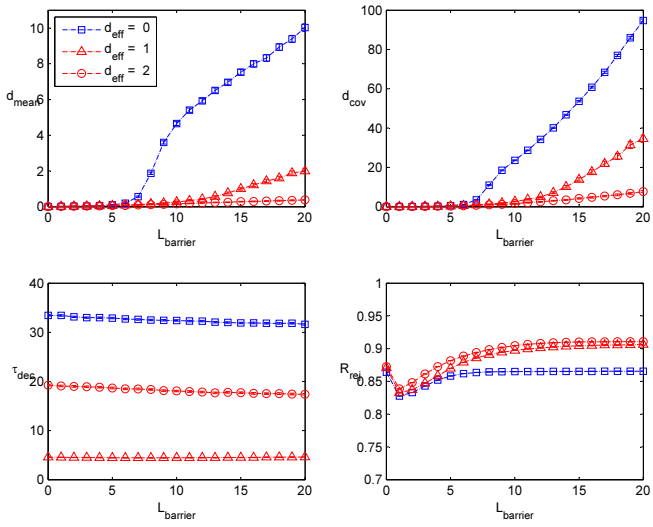

FIG. 7: Plots of the 2D free energy barrier tests.

[1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of State Calculations by Fast Computing Machines," *J. Chem. Phys.* **21** (1953) 1087–1092.

[2] W. K. Hastings, "Monte Carlo Sampling Methods Using Markov Chains and their Applications," *Biometrika* **57** no. 1, (1970) 97–109.

[3] J. J. Heckman, "Statistical Inference and String Theory," *Int. J. Mod. Phys.* **A30** no. 26, (2015) 1550160, `arXiv:1305.3621 [hep-th]`.

[4] R. H. Swendsen and J.-S. Wang, "Replica Monte Carlo Simulation of Spin-Glasses," *Phys. Rev. Lett.* **57** no. 21, (1986) 2607–2609.

[5] C. J. Geyer, "Markov Chain Monte Carlo Maximum Likelihood," in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, E. M. Keramidas, ed., pp. 156–163. Interface Foundation, 1991.

[6] W. R. Gilks, G. O. Roberts, and E. I. George, "Adaptive Direction Sampling," *Journal of the Royal Statistical Society. Series D (The Statistician)* **43** no. 1, (1997) 179–189.

[7] D. J. Earl and M. W. Deem, "Parallel tempering: Theory, applications, and new perspectives," *Phys. Chem. Chem. Phys.* **7** (2005) 3910–3916.

[8] R. M. Neal, "MCMC Using Ensembles of States for Problems with Fast and Slow Variables such as Gaussian Process Regression," `arXiv:1101.0387 [stat]`.

[9] J. Goodman and J. Weare, "Ensemble Samplers with Affine Invariance," *Comm. in Appl. Math. and Comp. Sci.* **5** no. 1, (2010) 65–80.

[10] J. J. Heckman, J. G. Bernstein, and B. Vigoda, "MCMC with Strings and Branes: The Suburban Algorithm (Extended Version)," `arXiv:1605.05334 [physics.comp-ph]`.

[11] S. Hershey, J. Bernstein, B. Bradley, A. Schweitzer, N. Stein, T. Weber, and B. Vigoda, "Accelerating Inference: towards a full Language, Compiler and Hardware Stack," `arXiv:1212.2991 [cs.SE]`.

[12] R. A. Neal, "Slice sampling," *The Ann. of Stat.* **31** no. 3, (2003) 705–767.

[13] A. Gelman, W. R. Gilks, and G. O. Roberts, "Weak convergence and optimal scaling of random walk Metropolis algorithms," *Ann. Appl. Prob.* **7** no. 1, (1997) 110–120.

[14] Y. Chen, M. Welling, and A. J. Smola, "Super-Samples from Kernel Herding," `arXiv:1203.3472 [cs.LG]`.

[15] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory.* Addison-Wesley, Reading, USA, 1995.

[16] A. Wipf, "Statistical approach to quantum field theory," *Lect. Notes Phys.* **864**, (2013).

[17] V. Balasubramanian, "Statistical Inference, Occam's Razor and Statistical Mechanics on the Space of Probability Distributions," *Neural Comp.* **9(2)** (1997) 349–368, `arXiv:cond-mat/9601030`.

[18] P. Di Francesco, P. Mathieu, and D. Senechal, *Conformal Field Theory.* Springer-Verlag, New York, USA, 1997.

Appendix A: Path Integrals for MCMC

In this Appendix we give a path integral formulation for MCMC with extended objects. For additional background on path integrals in statistical field theory, see [15, 16]. In what follows, we denote the random variable as X with outcome x on a target space Ω with measure dx. We consider sampling from a probability density $\pi(x)$. In accord with physical intuition, we view $-\log \pi(x) = V(x)$ as a potential energy.

In general, our aim is to discover the structure of $\pi(x)$ by using some sampling algorithm to produce a sequence of values $x^{(1)}, ..., x^{(N)}$. A quantity of interest is the expected value of $\pi(x)$ with respect to a given probability distribution of paths. This helps in telling us the relative speed of convergence and the mixing rate. To study this, it is helpful to evaluate the expectation value of the quantity:

$$\prod_{i=1}^{N} \exp(-V(x^{(i)})) \tag{A1}$$

with respect to a given path generated by our sampler.

In more general terms, the reason to be interested in this expectation value comes from the statistical mechanical interpretation of statistical inference [3, 17]: There is a competition between staying in high likelihood regions (minimizing the potential), and exploring more of the distribution (maximizing entropy). The tradeoff between the two is neatly captured by the path integral formalism: It tells us about a particle moving in a potential $V(x)$, and subject to a thermal background, as specified by the choice of probability measure over possible paths. Indeed, we will view this probability measure as defining a "kinetic energy" in the sense that at each time step, we apply a random kick to the trajectory of the particle, as dictated by its contact with a thermal reservoir.

Along these lines, if we have an MCMC sampler with transition probabilities $T(x^{(i)} \to x^{(i+1)})$, marginalizing over the intermediate values yields the expected value of line (A1):

$$\mathcal{Z} = \int [dx^{(t)}] \left(\prod_{i=0}^{N-1} T(x^{(i)} \to x^{(i+1)}) e^{-V(x^{(i+1)})} \right) \tag{A2}$$

where we have introduced the measure factor $[dx^{(t)}] = dx^{(1)}...dx^{(N)}$. We would like to interpret $V(x)$ as the potential energy and $-\log T(x^{(i)} \to x^{(i+1)})$ as a kinetic energy:

$$V = -\log \pi \quad \text{and} \quad K = -\log T, \tag{A3}$$

We now observe that our expectation value has the form of a well-known object in physics:

$$\mathcal{Z}(x^{\text{begin}} \to x^{\text{end}}) = \int_{\text{begin}}^{\text{end}} [dx^{(t)}] e^{-\sum_t L^{(E)}[x^{(t)}]}, \tag{A4}$$

A path integral! Here the Euclidean signature Lagrangian is:

$$L^{(E)}[x^{(t)}] = K + V. \tag{A5}$$

Since we shall also be taking the number of timesteps to be very large, we make the Riemann sum approximation and introduce the rescaled Lagrangian density:

$$\frac{1}{N} \sum_t \mapsto \int dt, \quad N L^{(E)} \mapsto \mathcal{L}^{(E)} \tag{A6}$$

so that we can write our process as:

$$\mathcal{Z}(x^{\text{begin}} \to x^{\text{end}}) = \int [dx] e^{-\int dt \mathcal{L}^{(E)}[x(t)]}, \tag{A7}$$

where by abuse of notation, we use the same variable t to reference both the discretized timestep as well as its continuum counterpart.

To give further justification for this terminology, consider now the specific case of the Metropolis-Hastings algorithm. In this case, we have a proposal kernel $q(x'|x)$, and acceptance probability:

$$a(x'|x) = \min\left(1, \frac{q(x|x')}{q(x'|x)} \frac{\pi(x')}{\pi(x)}\right). \tag{A8}$$

The total transmission probability is then given by a sum of two terms. One is given by $a(x'|x)q(x'|x)$, i.e., we accept the new sample. We also sometimes reject the sample, i.e., we keep the same value as before:

$$T(x \to x') = r \times \delta(x - x') + a(x'|x)q(x'|x), \tag{A9}$$

where $\delta(x - x')$ is the Dirac delta function, and we have introduced an averaged rejection rate:

$$r \equiv 1 - \int dx' \; a(x'|x)q(x'|x). \tag{A10}$$

To gain further insight, we now approximate the mixture model $T(x \to x')$ by a normal distribution $q_{\text{eff}}\left(x^{(t+1)}|x^{(t)}\right)$ such that $-\log q_{\text{eff}}\left(x^{(t+1)}|x^{(t)}\right) \sim \alpha_{\text{eff}}\left(x^{(t+1)} - x^{(t)}\right)^2$. Hence,[2]

$$L^{(E)}[x^{(t)}] \simeq \alpha_{\text{eff}}\left(x^{(t+1)} - x^{(t)}\right)^2 + V(x^{(t)}) + ..., \tag{A11}$$

where here, the "..." denotes additional correction terms which are typically suppressed by powers of $1/N$.

Our plan will be to assume a kinetic term with quadratic time derivatives, but a general potential. The overall strength of the kinetic term will depend on details

[2] For example, for a Gaussian proposal kernel with $-\log q(x^{(t+1)}|x^{(t)}) \sim \alpha(x^{(t+1)} - x^{(t)})^2$, matching the first and second moments to q_{eff} requires $\alpha_{\text{eff}} = \alpha/\overline{a}$, with $\overline{a}$ the average acceptance rate.

such as the average acceptance rate. As the acceptance rate decreases, α_{eff} increases and the sampled values all concentrate together.

We now turn to the generalization of the above concepts for strings and branes, i.e., extended objects. Introduce M copies of the original distribution, and consider the related joint distribution:

$$\pi(x_1, ..., x_M) = \pi(x_1)...\pi(x_M). \tag{A12}$$

If we keep the proposal kernel unchanged, we can simply describe the evolution of M independent point particles exploring an enlarged target space:

$$\Omega_{\text{enlarged}} = \Omega^M = \underbrace{\Omega \times ... \times \Omega}_{M}. \tag{A13}$$

If we also view the individual statistical agents on the worldvolume as indistinguishable, we can also consider quotienting by the symmetric group on M letters, S_M:

$$\Omega^{\mathcal{S}}_{\text{enlarged}} = X^M / S_M. \tag{A14}$$

Of course, we are also free to consider a more general proposal kernel in which we correlate these values. Viewed in this way, an extended object is a single point particle, but on an enlarged target space. The precise way in which we correlate entries across a grid will in turn dictate the type of extended object.

Indeed, much of the path integral formalism carries over unchanged. The only difference is that now, we must also keep track of the spatial extent of our object. So, we again introduce a potential energy V and a kinetic energy K:

$$V = -\log \pi \quad \text{and} \quad K = -\log T, \tag{A15}$$

and a Euclidean signature Lagrangian density:

$$L^{(E)}[x(t, \sigma_A)] = K + V, \tag{A16}$$

where here, σ_A indexes locations on the extended object, and the subscript A makes implicit reference to the adjacency on the graph. In a similar notation, the expected value is now:

$$\mathcal{Z}(x_{\text{begin}} \to x_{\text{end}}|A) = \int [dx]\ e^{-\sum_t \sum_\sigma L^{(E)}[x(t,\sigma_A)]}. \tag{A17}$$

Since we shall also be taking the number of time steps and agents to be large, we again make the Riemann sum approximation:

$$\frac{1}{N}\sum_t \mapsto \int dt, \quad \frac{1}{M}\sum_\sigma \mapsto \int d\sigma_A \quad NML^{(E)} \mapsto \mathcal{L}^{(E)} \tag{A18}$$

so that:

$$\mathcal{Z}(x^{\text{begin}} \to x^{\text{end}}|A) = \int [dx] e^{-\int dt d\sigma_A\ \mathcal{L}^{(E)}[x(t,\sigma_A)]}, \tag{A19}$$

in the obvious notation.

So far, we have held fixed a particular adjacency matrix. This is somewhat arbitrary, and physical considerations suggest a natural generalization where we sum over a statistical ensemble of choices. One can loosely refer to this splitting and joining of connectivity as "incorporating gravity" into the dynamics of the extended object, because it can change the notion of which statistical agents are nearest neighbors.[3] Along these lines, we incorporate an ensemble $\mathcal{A}$ of possible adjacency matrices, with some prescribed probability to draw a given adjacency matrix. The topology of an extended object dictates a choice of statistical ensemble $\mathcal{A}$.

Since we evolve forward in discretized time steps, we can in principle have a sequence of such matrices $A^{(1)}, ..., A^{(N)}$, one for each timestep. For each draw of an adjacency matrix, the notion of nearest neighbor will change, which we denote by writing $\sigma_{A(t)}$, that is, we make implicit reference to the connectivity of nearest neighbors. Marginalizing over the choice of adjacency matrix, we get:

$$\mathcal{Z}(x_{\text{begin}} \to x_{\text{end}}) = \int [dx][dA]\ e^{-\sum_t \sum_\sigma L^{(E)}[x(t,\sigma_{A(t)})]}, \tag{A20}$$

where now the integral involves summing over multiple ensembles: the spatial and temporal values with measure factor $dx^{(t)}_\sigma$, as well as the choice of a random matrix from the ensemble $dA^{(t)}$ (one such integral for each timestep). At a very general level, one can view the adjacency matrix as adding additional auxiliary random variables to the process. So in this sense, it is simply part of the definition of the proposal kernel.

1. Dimensions and Correlations

Following some of the general considerations outlined in reference [3], we now discuss the extent to which the extended nature of such objects plays a role in statistical inference and in particular MCMC.

To keep our discussion from becoming overly general, we specialize to the case of a hypercubic lattice of agents in d spatial dimensions arranged on a torus, and we denote a location on the grid by a d-component vector σ. We can allow for the possibility of a fluctuating worldvolume by making the crude substitution $d \mapsto d_{\text{eff}}$.

Consider the Gaussian proposal kernel of line (11). In a large lattice, we approximate the finite differences in one of the d spatial directions by derivatives of continuous functions. Expanding in this limit, various cross-terms

[3] It is not quite gravity in the worldvolume theory, because there is a priori no guarantee that our sum over different graph topologies will have a smooth semi-classical limit. Nevertheless, summing over different ways to connect the statistical agents conveys the main point that the proximity of any two agents can change.

cancel and we get for the proposal kernel:

$$\propto \exp\left(-2\beta\sum_{\sigma}\left((D_t x_\sigma)^2 + \sum_{k=1}^{d}(D_k x_\sigma)^2\right)\right). \quad \text{(A21)}$$

where D_k denotes a finite difference in the k^{th} spatial component of the d-dimensional lattice.

Just as in the case of the point particle, the transition rate defines a kinetic energy quadratic in derivatives (to leading order), with an effective strength dictated by the overall acceptance rate.

One of the things we would most like to understand is the extent to which an extended object can explore the hills and valleys of V. We perform a perturbative analysis, at first viewing V as a small correction to the Lagrangian. Starting from some fixed position x_*, consider the expansion of V around this point:

$$V(x) = V(x_*) + V'(x_*)(x - x_*) + \frac{V''(x_*)}{2}(x - x_*)^2 + ..., \quad \text{(A22)}$$

Each of the derivatives of $V(x)$ reveals another characteristic feature length of $V(x)$. These feature lengths are specified by the values of the moments for the distribution $\pi(x)$.

When $V = 0$, there is a well-known behavior of correlation functions which is given by eqn. (7).[4] There is thus a rather sharp change in the inferential powers of an extended object above and below $d_{\text{eff}} \sim 1$.

To understand the impact of a non-trivial potential, we introduce the notion of a "scaling dimension" for $x(t, \sigma)$ and its derivatives. This is a well-known notion, see [18] for a review. Just as we assign a notion of proximity in space and time to agents on a grid, we can also ask how rescaling all distances on the grid via:

$$N \mapsto \lambda N \;\; M \mapsto \lambda^d M \quad \text{(A23)}$$

impacts the structure of our continuum theory Lagrangian. The key point is that provided N and M have been taken sufficiently large, or alternatively we take λ sufficiently large, we do not expect there to be any impact on the physical interpretation.

Unpacking this statement naturally leads us to the notion of a scaling dimension for $x(t, \sigma)$ itself. Observe that rescaling the number of samples and number of agents in line (A23) can be interpreted equivalently as holding fixed N and M, but rescaling t and σ:

$$(t, \sigma) \mapsto (\lambda t, \lambda\sigma). \quad \text{(A24)}$$

Now, for our kinetic term to remain invariant, we need to *also* rescale $x(t, \sigma)$:

$$x(t, \sigma) \mapsto \lambda^{-\Delta} x(\lambda t, \lambda\sigma). \quad \text{(A25)}$$

[4] One way to obtain this scaling relation is to observe that the Fourier transform of $1/k^2$ in $d + 1$ dimensions exhibits the requisite power law behavior.

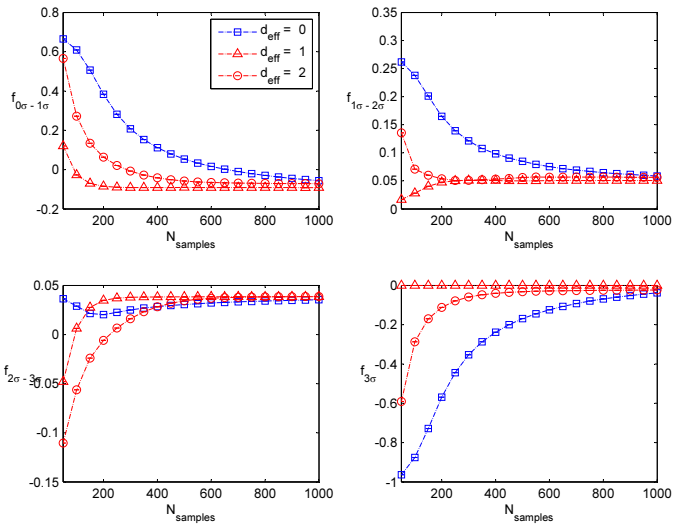

FIG. 8: Plots of the convergence to the true tail statistics for suburban samplers of the random seed 40 landscape.

The exponent Δ is often referred to as the "scaling dimension" for x obtained from "naive dimensional analysis" or NDA. It is "naive" in the sense that when the potential $V \neq 0$ and we have strong coupling, the notion of a scaling dimension may only emerge at sufficiently long distance scales. Note that because we are uniformly rescaling the spatial and temporal pieces of the grid, we get the same answer for the scaling dimension if we consider spatial derivatives along the grid. This assumption can also be relaxed in more general physical systems. To illustrate, invariance of the free field action requires:

$$\Delta = \frac{d - 1}{2}. \quad \text{(A26)}$$

We can also consider the behavior of a perturbation of the form $(x)^\mu (Dx)^\nu$. Applying our NDA analysis prescription, we see that under a rescaling, the contribution such a term makes to the action is:

$$\int dt d^d\sigma \; (x)^\mu (Dx)^\nu \mapsto \lambda^{-\mu\Delta - \nu(\Delta+1)+d+1} \int dt d^d\sigma \; (Dx)^2, \quad \text{(A27)}$$

so terms of the form $(Dx)^\nu$ for $\nu > 2$ die off as we take $N \to \infty$, i.e., $\lambda \to \infty$. Additionally, we see that when $d \leq 1$, we can in principle expect more general contributions of the form $(x)^\mu (Dx)^\nu$. For additional discussion on the interpretation of such contributions, see reference [3].

Consider next possible perturbations to the potential energy. Each successive interaction term in the potential is of the form x^n, with scaling dimension $n(d-1)/2$. So, for $d \leq 1$, all higher order terms can impact the long distance behavior of the correlation functions, while for $d > 1$, the most relevant term is bounded above, and the global structure of the potential will be missed.

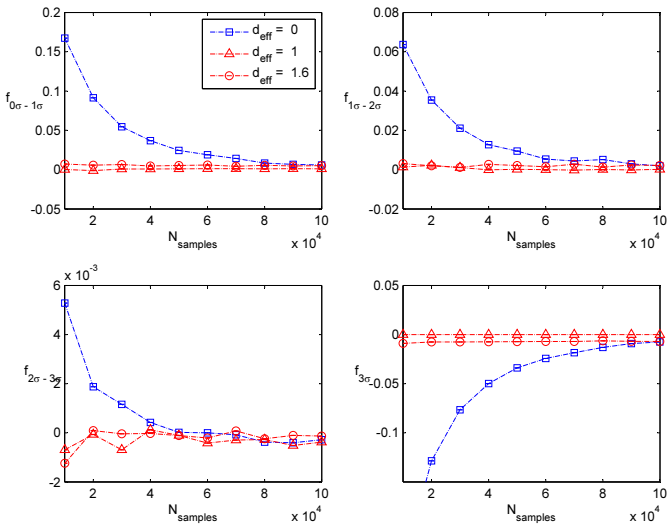

FIG. 9: Plots of the convergence to the true correct tail statistics for suburban samplers of the banana distribution.

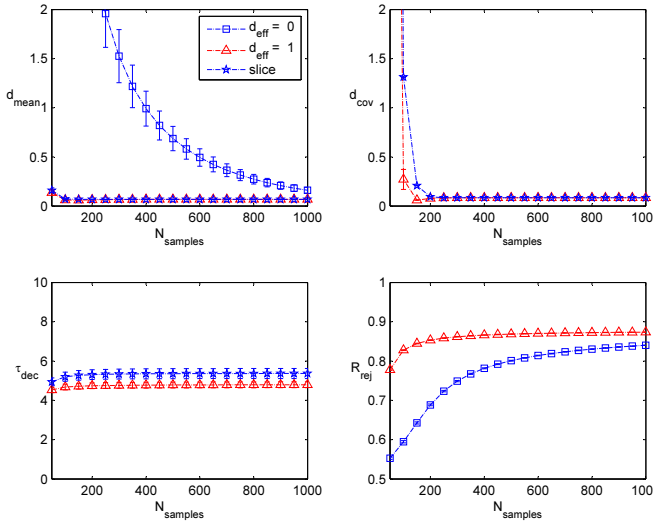

FIG. 10: Comparison between slice and suburban for the random seed 40 landscape.

Appendix B: Tail Statistics Tests

In figs. 8 and 9 we display some tests of how well a sampler collects "rare events," i.e., tail statistics. After taking a burn-in cut with $N_{\rm total}$ remaining samples, we compute the number of counts in the $0\sigma - 1\sigma$ region, the $1\sigma - 2\sigma$ region, the $2\sigma - 3\sigma$ region, and events which fall outside the 3σ region. For each such region, we compute the difference between the inferred and true counts and return the fraction:

$$f_{\rm region} \equiv \frac{N_{\rm inf} - N_{\rm true}}{N_{\rm total}}. \tag{B1}$$

Appendix C: Comparison with Slice

Figs. 10 and 11 compare the performance of a suburban sampler with $2d$ grid topology (with $d_{\rm eff} = 1$ and $\beta = 0.01$) with parallel slice within Gibbs sampling. We use the default implementation in `Dimple` so that for a 1D target distribution the initial size of the x-axis width is an interval of length one containing x_*, and the maximum number of doublings is 10. A direct comparison with suburban is subtle because in slice sampling the halting of the "stepping out" and "stepping in" loops is not fixed ahead of time. In practice we find that for a fixed number of samples, slice typically makes several more queries to the target distribution compared with suburban, roughly a factor of $\sim 5 - 10$. For large free energy barriers, it is also sometimes helpful to enlarge the initialization width from 1 to 100 (see fig. 11).

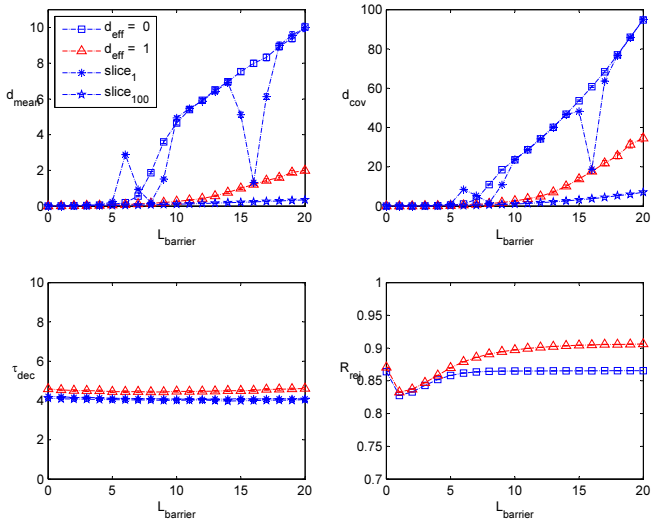

FIG. 11: Comparison between slice and suburban for the free energy barrier test. We show two different initialization widths for the slice sampler.